



2000

Computerworld
Smithsonian
Awards

A Search For New Heroes

CLICK FOR PRESS RELEASES

Congratulation to the
Class of 2000

HOME

**Long Summary****Benefits****Importance****Originality****Success****Difficulty****Naval Research Laboratory****LONG SUMMARY**

2000620 038

Please describe your application and the information technology used in conjunction with it. Please keep your language simple and your explanations non-technical.

Computer models have formed the basis for weather forecasting for more than 35 years. Global atmospheric models are run daily on the largest available supercomputers, dedicated to this single task, by national governments and super-national consortia. First the previous forecast and all available observations are blended in a data assimilation phase to form a "nowcast" of the current state of the atmosphere, and then this state is advanced in time to form a 3-10 day forecast. This must all be performed under real time constraints, typically within a 4-6 hour window. The results are then used directly for general weather forecasting and as input to higher resolution local models (e.g. for storm forecasting). Ocean forecasting is in principle similar to atmospheric forecasting, but with two major complications: (a) ocean eddies, at about 60 miles across, are typically 20 to 30 times smaller than comparable atmospheric highs and lows which means that roughly four orders of magnitude more computer time and three orders of magnitude more computer memory are required; and (b) there are relatively few observations below the ocean surface so data assimilation is effectively confined to using satellite observations

REPORT DOCUMENTATION PAGE			<i>Form Approved OMB No. 0704-0188</i>
<p>Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.</p>			
1. Agency Use Only (Leave Blank).	2. REPORT DATE June 2000	3. REPORT TYPE AND DATES COVERED Publication	
4. TITLE AND SUBTITLE 2000 Innovation Collection Case Study – Smithsonian's Permanent Research Collection		5. FUNDING NUMBERS Job Number Program Element 0601153N Project No. Accession No.	
6. AUTHOR(S) Harley Hurlburt			
7. PERFORMING ORGANIZATION AND ADDRESS Naval Research Laboratory Oceanography Division Stennis Space Center, MS 39529-5004		8. PERFORMING ORGANIZATION REPORT NUMBER NRL/PU/7320-399-396	
9. SPONSORING/MONITORING AGENCY NAMES(S) AND ADDRESS(ES) Office of Naval Research 800 North Quincy Street Arlington, VA 22217-5000		10. SPONSORING/MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES 2000 Computerworld Smithsonian Awards Electronic publication at web site http://innovate.si.edu , 7 pgs.			
12a. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release; distribution is unlimited		12b. DISTRIBUTION CODE	
<p>Computer models have formed the basis for weather forecasting for more than 35 years. Global atmospheric models are run daily on the largest available supercomputers, dedicated to this single task, by national governments and super-national consortia. First the previous forecast and all available observations are blended in a data assimilation phase to form "nowcast" of the current state of the atmosphere, and then this state is advanced in time to form a 3-10 day forecast. This must all be performed under real time constraints, typically within a 4-6 hour window. The results are then used directly for general weather forecasting and as input to higher resolution local models (e.g. for storm forecasting). Ocean forecasting is in principle similar to atmospheric forecasting, but with two major complications: (a) ocean eddies, at about 60 miles across, are typically 20 to 30 times smaller than comparable atmospheric highs and lows which means that roughly four orders of magnitude more computer time and three orders of magnitude more computer memory are required; and (b) there are relatively few observations below the ocean surface so data assimilation is effectively confined to using satellite observations of the surface. The duration of forecast skill for the ocean is not restricted to the 10 to 14 day limit for atmospheric highs and lows. We have demonstrated at least 30-day predictive skill for ocean eddies and the meandering of ocean currents and fronts, given sufficient ocean model resolution and satellite altimeter data from ERS-2 and TOPEX/POSEIDON. Forecasts of several months, or even years, are possible for some ocean phenomena.</p>			
14. SUBJECT TERMS TOPEX/POSEIDON, Smithsonian awards, global atmospheric modeling, and super computing			15. NUMBER OF PAGES 7
			16. Price Code
17. Security Classification of Report. Unclassified	18. Security Classification of This Page. Unclassified	19. Security Classification of Abstract Unclassified	20. Limitation of Abstract SAR

of the surface. The duration of forecast skill for the ocean is not restricted to the 10 to 14 day limit for atmospheric highs and lows. We have demonstrated at least 30-day predictive skill for ocean eddies and the meandering of ocean currents and fronts, given sufficient ocean model resolution and satellite altimeter data from ERS-2 and TOPEX/POSEIDON. Forecasts of several months, or even years, are possible for some ocean phenomena.

Ocean models divide the earth's surface into rectangular cells, discarding or ignoring the cells over land and solving "finite difference" approximations to continuous partial differential equations over the ocean cells. The continuous equations are in turn approximations to how we think the real ocean behaves. A major component of NRL's ocean modeling program has been a detailed study of how small the cells need to be for ocean prediction. We have strong evidence that an ocean model of this type (which includes all popular basin-scale ocean models) needs to use cells that are at most about 4 miles across at mid-latitudes. Doubling the resolution to 2 miles per cell gives substantial improvement but doubling again to 1 mile gives only modest additional improvement. This is for the global and basin-scale. Coastal models would use the global forecast for boundary conditions and would require much smaller cells, but covering only a limited region. At 2 miles the optimal resolution is finer than might be expected based on the size of eddies. In relation to ocean eddy size it is similar to the resolution currently used by the leading weather forecasting models in relation to the size of atmospheric highs and lows. More specifically, our research has shown that fine resolution of the ocean eddy scale is required to obtain coupling between upper ocean currents and seafloor topography, which can strongly affect the pathways of upper ocean currents and fronts, including the Gulf Stream in the Atlantic and the Kuroshio in the Pacific. This coupling occurs via turbulent flow instabilities in the upper ocean which generate deep ocean currents. These currents are guided by the seafloor topography. In turn the deep currents and thus the bottom have a steering effect on the upper ocean currents. The high resolution is also required to obtain sharp fronts that span major ocean basins and a nonlinear effect on the large scale C-shape of ocean gyres such as the Sargasso Sea in the Atlantic.

Depicting the evolution of ocean eddies and the meandering of ocean currents and fronts via data assimilation into the ocean model is a major challenge. Only satellite data has the potential for sufficient resolution with global coverage and then only at the ocean surface. However, NRL has developed statistical techniques for projecting this data downward that are effective even for deep currents which are only weakly correlated with surface currents. They are particularly effective when used in conjunction with the predictive skill of a high resolution ocean model and particularly effective for sea surface elevation data. Ocean currents in conjunction with temperature and salinity variations cause the ocean surface height to vary by about 6 feet globally, with a 1-3 foot change across eddies and fronts. This relatively small signal can be measured by satellite altimeters such as carried on ERS-2 and TOPEX/POSEIDON, but so far it is only measured along ground tracks directly beneath the satellite orbit. For TOPEX/POSEIDON, these have a 10-day repeat cycle with 189 mile equatorial track spacing and for ERS-2 a 35-day repeat and 48 mile equatorial track spacing. Hence another challenge for the data-assimilative ocean model is to fill in the space-time gaps between the altimeter tracks, again requiring an ocean model with predictive skill.

As far back as 1989, the President's Office of Science and Technology recognized global ocean modeling and prediction as a "Grand Challenge" problem, defined as requiring a computer system capable of sustaining at least one trillion floating point adds or multiplies per second. By taking a multi-pronged approach to cost minimization, we are solving the problem on systems capable of only a few percent of this performance. One prong is the use of experiment sequences that use the largest cell size possible and an ocean basin rather than the entire globe whenever possible. This only gets us so far, since in the end there is no substitute for small cells and a global domain. Another prong is the use of the NRL Layered Ocean Model, NLOM, which has been specifically designed for eddy-resolving global ocean prediction. It is tens of times faster than other ocean models, in operations required for a given result, for this specific problem. NLOM's performance is in turn due to a range of design decisions, the most important of which is the use of isopycnal (density tracking) layers in the vertical rather than the more usual fixed depth cells. Density is the natural vertical coordinate system for the ocean, and it allows seven NLOM layers to replace 100 or more fixed levels. Another important advantage of this approach is that there is less need to increase the number of density tracking layers as the horizontal cell size is reduced. With NLOM halving the cell size requires about 8x as much computer power (4x from the number of cells plus 2x from the required smaller time step), but with most other ocean models the number of cells in the vertical should also be doubled requiring about 16x as much computer power. There are no "free lunches", so NLOM sacrifices some capabilities for performance. The most important of these is that, except for straits, NLOM cannot provide even minimal coverage of shallow coastal regions. Surprisingly, it has shown excellent simulation skill for sea surface temperature, an unexpected result for a seven-layer model. A third prong is a widely portable NLOM computer code that targets large scalable computers with high speed network connections. NLOM exhibits very good scalability (wall time speedup as more processors are used) on such systems. For example, we have routinely run NLOM on up to 1,152 Cray T3E processors at a sustained speed of about 100 billion useful floating point operations per second. A final prong of our efficiency drive is the use of an inexpensive data assimilation scheme backed by a statistical technique for relating surface data to subsurface fields. The statistics are from an atmospherically-forced 20-year inter-annual simulation of the same ocean model, an application that requires a model with high simulation skill. This can be

thought of as a technique for using pre-processing computations that do not have hard wall time constraints to increase the accuracy and reduce the computations of the real time data assimilation.

So far it has been possible to run NLOM in demonstration mode with 2 mile mid-latitude resolution globally (72°S - 65°N) and 1 mile resolution over the basin-scale subtropical Atlantic (9°N - 51°N), including the Caribbean and Gulf of Mexico. While, at present, these require greater computer resources than practical for an operational product, they do give information on the value added of increasing resolution and insight into model performance at 4 mile resolution.

An eddy-resolving global ocean prediction system using NLOM with 4-mile resolution is scheduled for transition to operational testing in 2001 and should be in routine operation by 2002. It is currently running globally in atmospherically-forced simulation mode; and in near real time over the Pacific Ocean north of 20°S with assimilation of satellite altimeter data from ERS-2 and TOPEX/POSEIDON. The system gives a real time view of the ocean down to the 50-100 mile scale of ocean eddies and the meandering of ocean currents and fronts, a view with unprecedented resolution and clarity. We have also demonstrated that the data-assimilative Pacific Ocean model with 4 mile resolution can increase the effectiveness of satellite altimeter data as an observing system for the ocean, as well as providing forecast skill for 30 days or more for many ocean features. Eddy-resolving global ocean modeling and prediction has numerous military and civilian applications as outlined under "Benefits" below. In addition, the NRL effort is participating in a multinational Global Ocean Data Assimilation Experiment, dubbed GODAE, which is designed to help justify a permanent global ocean observing system by demonstrating useful real time ocean products with a customer base.

Naval Research Laboratory ORIGINALITY

What are the exceptional aspects of your project? Is it original? How? Is it the first, the only, the best or the most effective application of its kind? How did the project evolve? What is its background?

Generally a bimodal approach has been used in the ocean modeling community. Either models that were as simple as possible have been used to study an ocean process in isolation or the models have been made as complete as possible. The latter are very expensive computationally and their primary compromise is resolution. >From the beginning, over 20 years ago, we have had the long-term goal of developing a global and basin-scale ocean prediction capability for the Navy that was eddy-resolving, and the associated scientific goal of increased understanding of the ocean circulation. The need to resolve the eddies prompted us to take a different approach where cost-effectiveness of design choices was an essential consideration, especially in the areas of vertical coordinate design, model numerics and computational design.

As computing power increased, close consideration was given to the trade-off between increasing model resolution and adding capabilities which increased computational cost. Our ocean model performance was optimized for the available computing power on Navy global and basin-scale applications. Our basic research work on increased understanding of the ocean circulation made critical contributions to our effort and the choices we made. The result is an ocean model which is tens of times more efficient in computer time per model year than competing designs for a given domain and horizontal resolution, and the highest resolution global ocean simulation to date by a wide margin (2 miles at mid-latitudes vs 11 miles outside NRL). A similar approach has been used in data assimilation, where again there has been a tendency in the community to use relatively costly techniques at the expense of resolution. In our approach the 4-mile resolution of the ocean model contributes greatly to the effectiveness of the data assimilation process. The ocean model with assimilation has resulted in the first eddy resolving forecasts over a major ocean basin, the Pacific north of 20°S , and demonstrated forecast skill for at least 30 days.

We began in the late 1970s with very little in place: grossly inadequate computing power, a grossly inadequate ocean observing system, almost no previous work on data assimilation into ocean models, insufficient understanding of ocean dynamics, and no existing effort in place. All of these are now in place. Early work included seminal publications on the Kelvin wave theory of El Niño onset in the ocean and on eddy-shedding dynamics in the Gulf of Mexico as well as early development of our present ocean model, which began in 1976. The model was implemented with a vectorized computer code from the outset, running first on the NRL TI/ASC and later with autotasking on the NRL Cray XMP, which were then state of the art supercomputers.

In the mid 1980s we published papers on the potential for ocean prediction and the role of satellite altimeter data which largely outlined the vision and approach for our current global ocean modeling and prediction capability. We also performed pioneering research on oceanic data assimilation. At the time we had great concern that the extreme sparsity of subsurface

data, not obtainable from satellites, was a potentially crippling obstacle. Hence, in the 1980s and early 1990s we focused substantial effort on developing techniques for downward projection of the abundant satellite surface data in order to constrain sparsely observed subsurface fields. Our published results showed that constraint of even the very deep (abyssal) currents was possible (despite weak correlation with surface data and the lack of direct observations), and that in fact such constraint of the abyssal circulation was needed for skillful prediction of the upper ocean currents. During this time the capability of NLOM was also expanded substantially.

Starting in 1990, the computing power available to us was greatly increased with the establishment of a supercomputer center at the Naval Oceanographic Office (co-located with NRL ocean modelers at Stennis Space Center, MS) and the acquisition of a Cray YMP. In 1994 this became one of four Major Shared Resource Centers, available to all Defense Department researchers, under the DoD High Performance Computing Modernization Program. The TMC CM5 at NRL was also a major resource and the first massively parallel computer used for NLOM. This prompted the development of the current scalable portable computer code for NLOM, that can run efficiently and interchangeably on a wide variety of computing platforms. The modeling capability of NLOM was also increased, particularly in the areas of thermodynamics and the addition of a mixed layer that provides and accurate sea surface temperature, a task completed in 1999.

The enormous increases in computing power available to us allowed eddy-resolving global and basin-scale modeling with increasing resolution and capability during the 1990s and many firsts in ocean model resolution. It also allowed the discovery of impacts that resolution has on the realism of ocean circulation simulations and increased understanding of the ocean circulation. Examples include the shape of large-scale ocean gyres, ocean fronts that span major ocean basins, the pathways of the Gulf Stream in the Atlantic and the Kuroshio current in the Pacific, upper ocean - topographic coupling via turbulent flow instabilities (discussed in the long summary), pathways of the global-scale 3-D ocean circulation, and ocean variability. We also discovered a decadal impact of the 1982-3 El Niño which made the cover of Nature in August 1994. This discovery was chronicled in numerous newspapers and science-oriented magazines including the New York Times and Science News. It was named one of the top 75 science news stories of 1994 by Discover Magazine. All of this work led to the current global ocean modeling and prediction capabilities outlined above.

Naval Research Laboratory BENEFITS

Has your project helped those it was designed to help? In your opinion, how has it affected them? What new advantage or opportunity does your project provide to people? Has your project fundamentally changed how tasks are performed? In your opinion, have you developed a technology that may lead to new ways of communicating and processing information? What change might unfold?

The NLOM system with 4 mile resolution gives a real time view of the global ocean down to the 50-100 mile scale of ocean eddies and the meandering of ocean currents and fronts, a view with unprecedented resolution and clarity, and demonstrated forecast skill for a month or more for many ocean features. The assimilation of satellite altimeter data into this system will make more effective use of near real time altimeter data from ERS-2, TOPEX/POSEIDON, and their follow-ons via the Altimeter Data Fusion Center at the Naval Oceanographic Office. Other data, such as sea surface temperature and sparse vertical profiles of temperature and salinity, will be assimilated as well.

Eddy-resolving global ocean modeling and prediction has numerous military and civilian applications such as assimilation and synthesis of global satellite surface data; high resolution ocean nowcasting and forecasting; optimum ship track routing; search and rescue; anti-submarine warfare and surveillance; tactical planning; high resolution boundary conditions for even higher resolution coastal models; sea surface temperature for long range weather prediction; inputs for shipboard environmental products; environmental simulation and synthetic environments; observing system simulations; ocean research; pollution and tracer tracking; inputs to water quality assessment; El Niño monitoring and prediction; fisheries; and ocean structure design such as deep-sea oil platforms. We are participating in a multinational experiment, GODAE, designed to help justify a permanent global ocean observing system by demonstrating useful real time ocean products with a customer base. GODAE is scheduled for 2003-2005 (2007 in the U.S.) with pilot programs in 2000-2002. NRL is represented on both the International and U.S. GODAE Scientific Steering Teams.

In atmospherically-forced simulation mode the ocean model has proven a valuable research tool in gaining a better understanding of the ocean circulation. In this capacity it has been run in a variety of global and basin-scale configurations at 1 mile to 35 mile resolution. In this mode it has been used to provide realistic simulated data for testing data assimilation techniques in an environment where the "truth" is accurately known and it has been used to provide statistics for downward projection of the much more abundant surface data. In addition, the simulated data have been used for observing system

simulation to help gauge observing system performance characteristics, space-time resolution requirements, and observing system effectiveness. For example, simulated data from NLOM have been used in observing system simulation studies for satellite altimetry, acoustic tomography, and a global array of 3000 drifting floats designed to measure profiles of temperature and salinity.

Naval Research Laboratory SUCCESS

Has your project achieved or exceeded its goals? Is it fully operational? How many people benefit from it? If possible, include an example of how the project has benefited a specific individual, enterprise or organization. Please include personal quotes from individuals who have directly benefited from your work. Describe future plans for the project.

It has been a long road from our initial simulations in the Gulf of Mexico during the late 1970s, but all the pieces of the puzzle are coming together and we are on schedule for global ocean forecasts, with 4 mile mid-latitude resolution, to become routine by 2002. This does not require any additional breakthroughs in computer power or forecasting technique. The expected increase in computer capabilities over the next two years are sufficient to reduce our existing forecast system design to routine practice.

A global nowcast based on a 18 mile mid-latitude resolution ocean model which assimilates satellite altimeter data has been running daily in real time since early in 1998, with the results published on the web. At this resolution NLOM does not resolve fronts and eddies, and an alternative purely statistical product, MODAS, using optimal interpolation and also developed at NRL provides nowcasts of similar or better quality. However even at 18 mile resolution the NLOM based product is more dynamically consistent than MODAS, and is particularly useful in the tropics.

More recently we started running a 4 mile mid-latitude resolution Pacific Ocean model north of 20°S in near real time (i.e. updated every few days). In hindcast studies that follow standard nowcast and forecast procedures, but using data from a previous time period, we compared this Pacific model with the 18 mile mid-latitude resolution global model. Both assimilated satellite altimeter data from ERS-2 and TOPEX/POSEIDON and then performed month-long forecasts started from the data assimilative nowcast states. The Pacific cases are the very first basin-wide demonstration of skillful nowcasts and forecasts of oceanic fronts and eddies for any ocean basin. They also demonstrate that altimeter data alone is sufficient to produce an accurate nowcast when a high resolution ocean model is in the loop to fill in the space-time gaps in the altimeter data. The global model was much less successful in this feasibility demonstration, as expected from the earlier discussion of resolution requirements. In addition, the nowcast/forecast results at 4 mile resolution were substantially better than expected for a system that uses satellite altimeter data as the only observing system.

The global results for sea surface temperature (SST) have also exceeded expectations, particularly for a model with only seven layers in the vertical. The embedded mixed layer in NLOM gives accurate SST based on accurate atmospheric forcing even with no assimilation of SST data. With climatological atmospheric forcing, global NLOM gives SSTs accurate to within 1/2°C (1°F) for the annual mean and 0.7°C for the seasonal cycle. NLOM was run 1979-98 with 6-12 hourly atmospheric forcing and no assimilation of SST data, and then compared to 337 year-long daily time series of observed SST around the world over the 1980-98 time frame. The median error was less than 1°C and the median correlation coefficient was about 0.9, again with no assimilation of SST data. These results indicate that NLOM SST is sufficiently accurate to be used as a platform for assimilation of SST data (which has gaps due to cloudiness) and for SST forecasting.

Naval Research Laboratory DIFFICULTY

What were the most important obstacles that had to be overcome in order for your work to be successful? Technical problems? Resources? Expertise? Organizational problems? Often the most innovative projects encounter the greatest resistance when they are originally proposed. If you had to fight for funding, it would be useful to include a summary of the objections you faced and how you overcame them.

There have been numerous major technical challenges, as outlined above, but the dominant technical difficulty has always been that computers are not fast enough. We have received outstanding support from the Navy and from the DoD High Performance Computing Modernization Program, in providing the substantial resources on the largest state of the art supercomputers needed to tackle this Grand Challenge problem.

Another major challenge was developing an approach that would work well and that would be much more efficient than the conventional brute force approaches, particularly when a convincing proof by demonstration was not possible for more than a decade (that came in 1999) and when many peers, sponsors, managers, and customers were very skeptical (others were very supportive) that a model with only a few layers and only assimilating satellite altimetry would be adequate for eddy resolving ocean forecasting. There was widespread doubt that a model with a few layers could produce forecast skill for ocean eddies and the meandering of ocean currents and fronts, and as recently as 1998 leaders in the field argued at international meetings that satellite altimetry did not have adequate space-time resolution for nowcasting and forecasting of such features. However, in 1999 the 6-layer Pacific NLOM with 4 mile resolution demonstrated nowcast skill and greater than one month forecast skill for these features. Note that an operational system would assimilate additional data types, such as SST and sparse subsurface temperature and salinity profiles, but satellite altimetry is the primary data source.

Resistance has also come from those who expect a first generation forecast product to meet all the Navy's needs in this area. There are limitations to NLOM that are significant from a Navy perspective, primarily the lack of any nowcast/forecast in coastal regions (although 4 mile resolution is not nearly sufficient in shallow water), and others that appear significant but in practice are less important. For example, a detailed vertical profile is not a model output field but must be obtained as a post-processing step. These limitations are a side effect of being able to run the system at all on available computer systems. Our current expectation is that, absent an unforeseen breakthrough in technique, a global eddy resolving prediction system without these limitations will not be available until about 2007. In fact, the appropriate configuration for such a third generation product (the 2nd generation being based on NLOM at 2 mile resolution) is not certain today, and substantial work is necessary to sort out which of several promising approaches is best, possibly delaying deployment beyond 2007.

Naval Research Laboratory

Application : Eddy-Resolving Global Ocean
Modeling and Prediction

Category: Science

Status: Finalist

Nominated By: Cray Inc.

Advanced algorithms and high-performance supercomputers enable the prediction of ocean behavior and set the stage for future forecasts that will be accurate for months, and perhaps years, in advance.



Naval Research Laboratory IMPORTANCE

How did information technology contribute to this project? Describe any new technologies used and/or cite innovative uses of existing technology. For example, did you find new ways to use existing technology to create new benefits for society? Or, did you define a problem and develop new technology to solve it? How quickly has your targeted audience of users embraced your innovation? Or, how rapidly do you predict they will? Does your work define new challenges for society? If so, please describe what you believe they may be.

Global ocean modeling and prediction is a "Grand Challenge" problem, so progress can only be made by using the highest performance computer systems at close to maximum efficiency. Over the past ten years three distinct types of machine have in turn been the fastest for ocean modeling. First the Cray C90 was dominant, then the Thinking Machines Corporation CM5,

and finally the Cray/SGI T3E. These three machines each required a distinct programming style: autotasking plus vectorization for the Cray C90; data parallel for the TMC CM5; and message passing for the Cray/SGI T3E. Fortunately the basic algorithms used by ocean models are suitable for all three machine types, but the rapid change in programming styles required two complete rewrites of the original Cray C90 computer code. One of the lessons learned from trying to simultaneously maintain the Cray C90 and TMC CM5 codes was that a single source code across all machines types was a highly desirable goal. Therefore the latest NLOM code supports all three programming styles in a single source code, with 80% of the source code common to all styles. We are typically one of the first ocean modeling groups to try out a new machine, and as "early adopters" we often have to deal with unstable system software and the other problems of new machine architectures. This is made worthwhile by the opportunity to use the fastest machines from the earliest possible date. A recent example of this is our use of 1,152 Cray T3E processors to run a near-global (72°S-65°N) ocean simulation with 2 mile mid-latitude resolution. Cray T3E's are the mainstays of many ocean modeling projects, but typically only 100-200 processors are used at one time. Demonstrating that ocean models, and the Cray T3E, can scale to 1,000 nodes was an important milestone. The software modifications required to get good performance on 1,152 nodes were relatively minor, because earlier optimizations developed for 100-node machines that scale less well than the T3E were also applicable to a 1000-node Cray T3E. This was the first ever global simulation at our ultimate 2-mile target resolution for ocean forecasting. For comparison, 11 miles is the highest resolution to date for a global ocean model outside of NRL. The difference between 11 miles and 2 miles represents at least a 150-fold increase in computer resources. We don't have 100 times faster computers than other ocean modeling groups, but instead we have taken aggressive steps to reduce the cost per equivalent result for ocean simulation and forecasting. For example, NLOM uses a semi-implicit time stepping scheme and solves the resulting Helmholtz's equation every time step using the Capacitance Matrix Technique, CMT. This is one of the most important algorithmic optimizations in NLOM because it allows a much longer time step. The CMT is very fast and direct (i.e. non-iterative) but it replaces a huge sparse system of linear equations over the ocean with a large dense (fully populated) system on the coastline. For the 2-mile global domain the coastline matrix is 83,000 by 83,000, and this may be the largest system of dense linear equations ever solved at every time step of a time dependent problem. Fifty gigabytes of memory is required just to hold these matrices, so the CMT can be thought of as a method that trades memory for performance. Since memory capacity is increasing faster than processor performance this is a very worthwhile trade.

It is still unclear which machine, or class of machines, will take over from the Cray T3E as the next "king of the ocean modeling hill". Low cost clusters of personal computers are becoming very popular, but for use in global ocean modeling a cluster would need 1,000 or more PCs. Even on small clusters, of perhaps 50 PCs, the existing low cost network interconnects are not fast enough for ocean models. One possibility for a lower cost alternative to traditional high performance computer systems is a cluster of multi-processor nodes. The advantage of multiple processors is that they can perform enough work to potentially hide almost all the communication between nodes, even on a relatively slow network. We are actively working on modifications to NLOM's source code to better support such systems. No matter what future machine type emerges as optimal for ocean modeling, NLOM will be one of the first ocean models taking full advantage of the new architecture.

SEARCH 2000 COLLECTION

For information on all members of the 2000 Collection,
search by Category, Company Name, Nominating
Company, Application, City, State, Country or
Keywords according to your area of interest.

Category

Laureate Name City

Nominating Company State

Application Country

Keywords

Search **Clear**